arXiv:1604.01502v1 [math.CO] 6 Apr 2016

Distinct and repeated distances on a surface and incidences between points and spheres *

Micha Sharir† Noam Solomon‡

April 7, 2016

Abstract

In their seminal paper from 2010, Guth and Katz [24] proved that the number of distinct distances determined by a set of n points in $\mathbb{R}^2$ is $\Omega(n/\log n)$, thus (almost) settling Erdős's distinct distances problem, open for nearly 65 years. In $\mathbb{R}^3$, it is conjectured that a set of n points determines at least $\Omega(n^{2/3})$ distinct distances. This bound is best possible as it is attained by the vertices of the $n^{1/3} \times n^{1/3} \times n^{1/3}$ integer grid. This problem however is still wide open, for many years. The best known lower bound is due to Solymosi and Vu [39].

In this paper we show that the number of distinct distances determined by a set of n points on a constant-degree two-dimensional algebraic variety V (i.e., a surface) in $\mathbb{R}^3$ is at least $\Omega\left(n^{7/9}/\text{polylog}\, n\right)$. This bound is significantly larger than the conjectured bound $\Omega(n^{2/3})$ for general point sets in $\mathbb{R}^3$.

We also show that the number of unit distances determined by n points on a surface V, as above, is $O(n^{4/3})$, a bound that matches the best known planar bound, and is worst-case tight in 3-space. This is in sharp contrast with the best known general bound $O(n^{3/2})$ for points in three dimensions.

We also obtain sharp bounds for bipartite versions of the distinct distances and the repeated distances problems.

To prove these results, we establish an improved upper bound for the number of incidences between a set P of m points and a set S of n spheres, of arbitrary radii, in $\mathbb{R}^3$, provided that the points lie on an algebraic surface V of constant degree, which does not have linear or spherical components. Specifically, the bound is

$$O\left(m^{2/3}n^{2/3} + m^{1/2}n^{7/8}\log^{\beta}(m^4/n) + m + n + \sum_c |P_c| \cdot |S_c|\right),$$

where the constant of proportionality and the constant exponent β depend on the degree of V, and where the sum ranges over all circles c that are fully contained in V, so that, for each such c, $P_c = P \cap c$ and S_c is the set of the spheres of S that contain c. In addition, $\sum_c |P_c| = O(m)$ and $\sum_c |S_c| = O(n)$.

This bound too improves upon earlier known bounds. These have been obtained for arbitrary point sets but only under severe restrictions about the spheres, which are dropped in our result.

Another interesting application of our result is an incidence bound for arbitrary points and spheres in 3-space, where we improve and generalize the previous work of Apfelbaum and Sharir [5].

Keywords. Combinatorial geometry, incidences, the polynomial method, algebraic geometry, distinct distances.

*Work on this paper by Noam Solomon and Micha Sharir was supported by Grant 892/13 from the Israel Science Foundation. Work by Micha Sharir was also supported by Grant 2012/229 from the U.S.–Israel Binational Science Foundation, by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

†School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. *michas@tau.ac.il*

‡School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. *noam.solom@gmail.com*

1 Introduction

Incidences between points and spheres. Let P be a set of m points, and S a set of n spheres of arbitrary radii in $\mathbb{R}^3$. Assume that P is contained in some two-dimensional algebraic variety (surface) V of constant degree D, which does not have any planar or spherical components. We wish to bound the size $I(P,S)$ of the *incidence graph* $G(P,S)$, whose edges connect all pairs $(p,s) \in P \times S$ such that p is incident to s. In general, and in this special setup too, $I(P,S)$ can be as large as the maximum possible value $|P| \cdot |S|$, by placing all the points of P on a circle, and make all the spheres of S contain this circle, in which case $G(P,S) = P \times S$. The bound that we are going to obtain will of course acknowledge this possibility, and will be of the form $I_0(P,S) + \sum_i |P_i| \cdot |S_i|$, where, for each i, $P_i \subseteq P$, $S_i \subseteq S$, and $P_i \times S_i \subseteq G(P,S)$. Moreover, each subgraph $P_i \times S_i$ is induced by a circle contained in V that contains all the points of P_i and is contained in all the spheres of S_i. Informally, the residue term $I_0(P,S)$ bounds the number of "accidental" incidences, those that cannot be "explained" in terms of large complete bipartite subgraphs of $G(P,S)$. The quality of the bound will be measured by $I_0(P,S)$ and by $\sum_i (|P_i| + |S_i|)$.

Distinct and repeated distances in $\mathbb{R}^3$. There are two main motivations for studying point-sphere incidences. One involves repeated and distinct distances. For repeated distances, one draws a sphere with the given distance as radius around each input point, and the number of incidences between the points and spheres is exactly twice the number of repetitions of that distance. Applications of this kind include [18, 29, 42]. For distinct distances, one draws spheres centered at the given points and having as radii all the t possible distances. The number of incidences of these nt spheres with the n given points is exactly $n(n-1)$, so an upper bound on point-sphere incidences will yield a lower bound on t; see, e.g., [1] and [40] for applications of this approach.

The point-sphere incidence approach is an effective tool for providing lower bounds on the number of distinct distances on a surface in three dimensions, as demonstrated below in Theorem 1.2.

A second, closely related class of applications involves the number of repetitions of more involved patterns, typically congruent and similar simplices in a given point set; see [2, 3, 12] for examples of such applications. These applications are not discussed in the present paper.

Background. Earlier works on point-sphere incidences have considered the general setup, where the points of P are arbitrarily placed in $\mathbb{R}^3$. Initial partial results go back to Chung [16] and to Clarkson et al. [18], and continue with the work of Aronov et al. [6]. Later, Agarwal et al. [2] have bounded the number of *non-degenerate* spheres with respect to a given point set, which was then improved by Apfelbaum and Sharir [5].[1] Recently, Zahl [42] gave a bound for the number of incidences between m points and n spheres in $\mathbb{R}^3$, when every triple of spheres intersect at a finite set of points (which is the general case), as part of a more general bound on the number of incidences between points and bounded-degree surfaces in $\mathbb{R}^3$ satisfying certain favorable conditions. Zahl's bound for spheres is $O(m^{3/4}n^{3/4} + m + n)$. (This bound was later generalized by Basu and Sombra [11] to incidences between points and bounded degree hypersurfaces in $\mathbb{R}^4$ satisfying certain analogous conditions.) The case of incidences with unit spheres have also been studied in Kaplan et al. [29], with the same upper bound. Other mildly related recent works include [10, 17, 33].

The study in this paper continues a similar recent study by the authors [37], involving incidences between points on a variety and planes in three dimensions.

[1] Given a finite point set $P \subset \mathbb{R}^3$ and a constant $0 < \eta < 1$, a sphere $\sigma \subset \mathbb{R}^3$ is called η-degenerate (with respect to P), if there exists a circle $c \subset \sigma$ such that $|c \cap P| \geq \eta |\sigma \cap P|$.

Main Theorem. As we show in this paper, the bound can be substantially improved when all the points of P lie on a constant-degree surface V. Our main result is the following theorem.

Theorem 1.1. *Let P be a set of m points on some algebraic surface V of constant degree D in $\mathbb{R}^3$, which has no linear or spherical components, and let S be a set of n spheres, of arbitrary radii, in $\mathbb{R}^3$. The incidence graph $G(P,S)$ can be decomposed as*

$$G(P,S) = G_0(P,S) \cup \bigcup_i (P_i \times S_i), \tag{1}$$

such that, for each i, $P_i \subseteq P$ and $S_i \subseteq S$, and

$$\begin{aligned} |G_0(P,S)| &= O\left(m^{2/3}n^{2/3} + m^{1/2}n^{7/8}\log^{\beta}(m^4/n) + m + n\right), \\ \sum_i |P_i| = O(m), \quad &\text{and} \quad \sum_i |S_i| = O(n), \end{aligned} \tag{2}$$

where the constant exponent β and the constants of proportionality depend on the degree D of V. Moreover, for each i there exists a circle $c_i \subset V$, such that $P_i = P \cap c_i$ and S_i is the set of spheres in S that contain c_i.

Remark. Apart from the improved bound on $|G_0(P,S)|$, our bound is stronger, when compared to earlier works, also in that it does not impose any restrictions on $G(P,S)$ (like not containing a complete bipartite graph of some fixed size), and gives a precise representation of graphs $G(P,S)$ that do have "too many" incidences. (An earlier attempt at characterizing such large graphs is given in Apfelbaum and Sharir [4] for the case of planes (and hyperplanes in higher dimensions). Although it caters to the general case (not requiring the points to lie on a surface), it is much weaker than our representation.)

The cases where V is (or contains) a plane or a sphere. We have explicitly ruled out these cases in our assumptions, because the situation in these cases is different. These cases are treated in Theorem 1.4 below. In these cases, each sphere intersects V in a circle, and the problem boils down to one involving incidences between points and circles in the plane, or on the sphere, except that the circles can have (potentially large) multiplicities. Problems of this sort (without multiplicity of the circles) have already been tackled in [1, 8, 30], and a suitable extension of the analysis in these papers can also handle multiplicities in a rather straightforward manner.

Applications. We begin by presenting two applications of Theorem 1.1. Actually, in the second application and in the first part of the first one, we get better bounds, because the spheres that arise there have a more constrained structure. First, we obtain the following lower bounds on the number of distinct distances involving points on a surface in $\mathbb{R}^3$.

Theorem 1.2. *(a) Let P be a set of n points on an algebraic surface V of constant degree D in $\mathbb{R}^3$, which has no linear or spherical components. Then the number of distinct distances determined by P is $\Omega(n^{7/9}/\log^{\beta_1} n)$, where the constant exponent β_1 and the constant of proportionality depend on the degree D of V.*

(b) Let P_1 be a set of m points on a surface V as in (a), and let P_2 be a set of n arbitrary points in $\mathbb{R}^3$. Then the number of distinct distances determined by pairs of points in $P_1 \times P_2$ is

$$\Omega\left(\min\left\{m^{4/7}n^{1/7}/\log^{\beta_2}(m^4/n),\ m,\ n\right\}\right),$$

where the constant exponent β_2 and the constant of proportionality depend on D.

While we believe that the bounds in the theorem are not tight, we note that the bound in (a) is significantly larger than the conjectured best-possible lower bound $\Omega(n^{2/3})$ for arbitrary point sets in $\mathbb{R}^3$, and so is the (somewhat weaker) bound in (b), for a suitable range of values m, n (including the case $m = n$).

As another application, we bound the number of unit (or repeated) distances involving points on a surface V, as above.

Theorem 1.3. *(a) Let P be a set of n points on some algebraic surface V of constant degree D in $\mathbb{R}^3$. Then P determines $O(n^{4/3})$ unit distances, where the constant of proportionality depends on the degree D of V.*

(b) Let P_1 be a set of m points on a surface V as in (a), and let P_2 be a set of n arbitrary points in $\mathbb{R}^3$. Then the number of unit distances determined by pairs of points in $P_1 \times P_2$ is

$$O\left(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\log^{\beta_3}(m^3/n) + m + n\right),$$

where the constant exponent β_3 and the constant of proportionality depend on D.

The bound in (a) matches the best known upper bound for points in the plane or on a sphere (see Brass, Moser, and Pach [14] for a review of known results), and is in fact worst-case tight, since a matching lower bound $\Omega(n^{4/3})$ is known for points on a sphere with radius $1/\sqrt{2}$. The bound in (b) (say, for $m = n$) is "in between" the best known general upper bound of $O(n^{3/2})$ for any set of n points in $\mathbb{R}^3$ (see [29, 42]) and the bound in (a).

Another interesting application of Theorem 1.1 is the following general point-sphere incidence bound in three dimensions. It improves the bound in Apfelbaum and Sharir [5], and is more general, since it does not assume the spheres to be non-degenerate, as is the case in [5]. This theorem is analogous to the works of Brass and Knauer [13] and Apfelbaum and Sharir [4], who studied incidences between points and planes (instead of spheres) in $\mathbb{R}^3$ (and hyperplanes in higher dimensions).

Theorem 1.4. *Let P be a set of m points in $\mathbb{R}^3$, and let S be a set of n spheres, of arbitrary radii, in $\mathbb{R}^3$. The incidence graph $G(P, S)$ can be decomposed as*

$$G(P, S) = G_0(P, S) \cup \bigcup_i (P_i \times S_i), \tag{3}$$

such that, for each i, $P_i \subseteq P$ and $S_i \subseteq S$, and

$$\begin{aligned} |G_0(P, S)| &= O\left(m^{8/11+\varepsilon}n^{9/11} + m + n\right), \\ \sum_i \left(|P_i| + |S_i|\right) &= O\left(m^{8/11+\varepsilon}n^{9/11} + m + n\right), \end{aligned} \tag{4}$$

for any $\varepsilon > 0$, where the constant of proportionality depends on ε. Moreover, for each i there exists a circle c_i, such that $P_i = P \cap c_i$ and S_i is the set of spheres in S that contain c_i.

The technique. Our approach continues the recent methodology of applying tools from algebra and algebraic geometry to problems in combinatorial (and computational) geometry, pioneered by Guth and Katz's works [23, 24]. The main tool in this methodology is the *polynomial partitioning technique*, which yields a divide-and-conquer mechanism via space decomposition, which in many instances is a more effective tool than more traditional space decomposition techniques (such as cuttings and simplicial partitions; see, e.g., [15]). Interestingly though, while we do use algebraic techniques, a major part of the analysis,

involving decomposition in a dual four-dimendsional space, goes back to the source, and applies a standard cutting-based decomposition, exploiting the fact that the objects that arise in this duality are points and hyperplanes. This requires a somewhat more careful analysis, but results in slightly improved bounds.

Theorem 1.1 will be proved in Section 2, Theorems 1.2 and 1.3 in Section 3, and Theorem 1.4 in Section 4.

2 Proof of Theorem 1.1

Let P, V, S, m, and n be as above. We first restrict the analysis to the case where V is irreducible. This can be done, without loss of generality, by decomposing V into its irreducible components, assign each point of P to each component that contains it, and assign the spheres of S to all the components. This decomposes the problem into at most D (as a matter of fact, at most $D/2$) subproblems, each involving an irreducible surface, and it thus follows that the original incidence count is at most $D/2$ times the bound for the irreducible case. In the remainder of this section we thus assume that V is irreducible.

To obtain the bound (2) on $I(P, S)$, we first derive a weaker bound, and then improve it via a suitable decomposition of dual space, similar to the way it has been done for circles in [1, 8], and in more generality in [38], and also resembles the handling of the simpler case of points and planes in the companion paper [37].

A basic weak bound. By Sharir, Sheffer, and Zahl [35, Lemma 3.2], except for at most two *"popular"* points, each point $p \in V$ is incident to at most $44D^2 = O(1)$ circles that are fully contained in V; this follows since V is neither a sphere nor a plane.

The number of incidences between the popular points, if any of them is in P, and the spheres of S is at most $2n$, so in what follows we ignore these points and assume that P does not contain a popular point.

Let C denote the set of circles that are fully contained in V, contain at least one point of P, and are contained in at least one sphere of S. For each circle $c \in C$ we form the bipartite subgraph $P_c \times S_c$ of $G(P, S)$, where $P_c = P \cap c$ and S_c is the set of all the spheres of S that contain c.

The preceding property therefore implies that $\sum_c |P_c| = O(m)$. As V is irreducible and non-spherical, it does not contain any of the spheres in S. Thus, for each $s \in S$, the intersection $s \cap V$ is an algebraic curve of degree at most $2D$ (as follows, e.g., from the generalized version of Bézout's theorem [21]), and can therefore contain at most $D = O(1)$ circles of C. This implies that $\sum_c |S_c| = O(n)$.

To recap, we have obtained a collection of complete bipartite graphs $P_c \times S_c$, so that $\bigcup_c (P_c \times S_c)$ is a portion of $G(P, S)$, $\sum_c |P_c| = O(m)$, and $\sum_c |S_c| = O(n)$.

An interesting special case is when V is *ruled* by circles. That is, each point $p \in V$ is incident to a circle that is fully contained in V. Actually, as follows from a generalization of the Cayley–Salmon theorem, established by Guth and Zahl [25], an irreducible surface of degree D that is not ruled by circles can fully contain at most cD^2 circles, for some absolute constant c. That is, if V is not ruled by circles, we also get a bound of $O(D^2) = O(1)$ on the *number* of complete bipartite graphs in the decomposition (1).

Surfaces ruled by circles contain infinitely many circles, but only finitely many of them will yield nonempty bipartite graphs $P_c \times S_c$. Informally, this means that we might get more complete bipartite subgraphs in $G(P, S)$, but each with a smaller number of edges; the linear bounds on the total size of their vertex sets continue to hold.

For each $s \in S$, put $\gamma_s := (s \cap V) \setminus \bigcup C$. As already observed, V does not fully contain any sphere of S, so each γ_s is at most one-dimensional. By construction, it does not contain any circle, and it might also be empty (for this or for other reasons). Note that if $s \cap V$ does contain a circle c, then $c \in C$, and the

incidences between s and the points of P on c are all already recorded in $P_c \times S_c$ (assuming, of course, that $P_c = c \cap P \neq \emptyset$; otherwise, removing c incurs no loss of incidences). Finally, we ignore the isolated points of γ_s. The number of such points on a sphere s is $O(1)$,[2] so the number of incidences (p, s), where p is an isolated point of γ_s, is at most $O(n)$. We thus obtain the decomposition in (1), by letting $G_0(P, S)$ denote the remaining portion of $G(P, S)$, after pruning away the complete bipartite graphs $P_c \times S_c$. We further remove from $G_0(P, S)$ all the $O(n)$ incidences involving isolated points on their incident spheres, continue to denote the resulting subgraph as $G_0(P, S)$, and put $I_0(P, S) = |G_0(P, S)|$.

Let Γ denote the set of the n curves γ_s, for $s \in S$. The curves of Γ are (spherical) algebraic curves of degree at most $2D$ (e.g., see [21]), and any pair of curves $\gamma_s, \gamma_{s'} \in \Gamma$ intersect in at most $2D = O(1)$ points. Indeed, any of these points is an intersection point of V with the circle $c = s \cap s'$; if c is fully contained in V then, by construction, it has been removed from both curves, and if c is not contained in V, it can intersect it in at most $2D$ points.

Note that $I_0(P, S)$ is equal to the number of incidences $I(P, \Gamma)$ between the points of P and the curves of Γ. To bound the latter quantity we proceed as follows.

We slightly tilt the coordinate frame to make it generic, and then project the curves of Γ onto the xy-plane. A suitable choice of the tilting guarantees that (i) no pair of intersection points or points of P project to the same point, (ii) if p is not incident to γ_s then the projections of p and of γ_s remain non-incident, and (iii) no pair of curves in Γ have overlapping projections. In addition, by construction, no curve of Γ contains any (vertical) segment. Let P^* and Γ^* denote, respectively, the set of projected points and the set of projected curves; the latter is a set of n plane algebraic curves of constant maximum degree $2D$ (see, e.g., Harris [27] for the fact that projections do not increase the degree). Moreover, $I(P, \Gamma)$ is equal to the number $I(P^*, \Gamma^*)$ of incidences between P^* and Γ^*.

By the recent result of Sharir and Zahl [38] (see Theorem 2.1 below), applied to Γ^*, the curves of Γ^* can be cut into $O(n^{3/2} \log^\kappa n)$ connected Jordan subarcs, where the constant exponent κ and the constant of proportionality depend on D, so that each pair of subarcs intersect at most once; the new arcs thus form a collection of *pseudo-segments*.

We can now apply the crossing-lemma technique of Székely [40], exactly as was done in Sharir and Zahl [38]. Since the resulting subarcs form a collection of pseudo-segments, and the number of their intersections is $O(n^2)$, Székely's analysis yields the bound

$$I_0(P, S) = I(P, \Gamma) = I(P^*, \Gamma^*) = O\left(m^{2/3}n^{2/3} + m + n^{3/2} \log^\kappa n\right).$$

Adding incidences recorded in the complete bipartite decomposition, as constructed above, and the $O(n)$ incidences with isolated points, we get our initial (weak) bound.

$$I(P, S) = O\left(m^{2/3}n^{2/3} + m + n^{3/2} \log^\kappa n + \sum_c |P_c| \cdot |S_c|\right), \tag{5}$$

where $\bigcup_c (P_c \times S_c)$ is contained in the incidence graph $G(P, S)$, $\sum_c |P_c| = O(m)$, and $\sum_c |S_c| = O(n)$.

The case $m = O(n^{1/4})$. Before proceeding to improve the bound in (5), we first dispose of the case $m = O(n^{1/4})$. As above, we first remove all the complete bipartite graphs $P_c \times S_c$, for $c \in C$, from $G(P, S)$. We then proceed to estimate $I_0(P, S)$ as follows. We call a sphere $s \in S$ *strongly degenerate* (or *degenerate*[3] for short) if all the points of $P \cap s$ are cocircular. We claim that the number of incidences between the points of P and the non-degenerate spheres is $O(m^4 + n) = O(n)$. Indeed, first discard the

[2] The projection of $s \cap V$ onto a generic plane is a plane algebraic curve of degree at most $2D$ [27], and isolated points are projected onto isolated points. By Harnack's curve theorem [26], the number of isolated points is thus $O(D^2) = O(1)$.

[3] This is much more restrictive than the notion of η-degeneracy mentioned earlier.

spheres $s \in S$ containing at most three points of P, losing at most $O(n)$ incidences. For an incidence between a point $p \in P$ and a surviving non-degenerate sphere $s \in S$, there exist (at least) three distinct points $q, q', q'' \in (P \setminus \{p\}) \cap s$ such that p, q, q', q'' do not all lie on a common circle; this follows since s is non-degenerate and $|s \cap P| \geq 4$, so s contains at least three points that are not all cocircular with p. The ordered quadruple (p, q, q', q'') therefore *uniquely* accounts for the incidence between p and s, and there are $O(m^4) = O(n)$ such quadruples.

To bound the number of incidences between the points of P and the degenerate spheres of S, fix a degenerate sphere $s \in S$, and assume that $m_s := |s \cap P| \geq 2D + 1$; the overall number of incidences on the other spheres is at most $2Dn = O(n)$. By assumption, all points of $s \cap P$ lie on a common circle c. Since c contains at least $2D + 1$ points of P, it must be contained in V, so the incidences involving s are all recorded in the complete bipartite graph $P_c \times S_c$. In other words, we have shown, for $m = O(n^{1/4})$,

$$I(P, S) = O\left(n + \sum_c |P_c| \cdot |S_c|\right), \tag{6}$$

where, as above, $\bigcup_c (P_c \times S_c)$ is contained in the incidence graph $G(P, S)$, and $\sum_c |P_c| = O(m)$ and $\sum_c |S_c| = O(n)$.

Remark. It seems likely that, with some care, the bound in (6) could also be obtained by the technique of Fox et al. [20, Corollary 2.3] (see also [38]).

Improving the bound. As in the analysis of incidences between points and circles (or pseudo-circles) in [1, 8], and the more general analysis in [38], the first two terms in (5) dominate when $m = \Omega(n^{5/4} \log^{3\kappa/2} n)$. When m is smaller, the third term, which is independent of m, is the one that dominates, and we then sharpen it as follows (a similar general approach is also used in [1, 8, 37, 38]).

We apply the standard lifting transform to 4-space, which maps each point (x, y, z) to the point $(x, y, z, x^2+y^2+z^2)$ on the paraboloid $w = x^2+y^2+z^2$, and maps each sphere $(x-a)^2+(y-b)^2+(z-c)^2 = r^2$ to the hyperplane $w = 2ax + 2by + 2cz + (r^2 - a^2 - b^2 - c^2)$. This lifting preserves incidences: a point is incident to a sphere iff the lifted point is incident to the lifted hyperplane. We next apply a standard duality in 4-space that maps points to hyperplanes and vice versa, and preserves point-hyperplane incidences. We denote by a^* the lifted and dualized image (hyperplane or point) of an object a (point or sphere); to simplify the terminology, we call a^* simply the *dual image* of a.

We thus get a set S^* of n points, and a set P^* of m hyperplanes in 4-space. We choose a parameter r, to be fixed later, and construct a $(1/r)$-*cutting* Ξ for P^* (see, e.g., [15]), which partitions $\mathbb{R}^4$ into $O(r^4)$ simplices, each crossed by at most m/r dual hyperplanes.

Incidences with boundary dual points. Let us first handle dual points that lie on the boundaries of the simplices of the $(1/r)$-cutting Ξ and the dual hyperplanes.

(i) Each dual point that lies in the relative interior of a 3-face φ of some simplex of Ξ has one incidence with the dual hyperplane that contains φ (if any), for a total of $O(n)$ such incidences. If such a point s^* is incident to another dual hyperplane p^*, then p^* must cross each simplex that has φ as a face (there are one or two such simplices). We then assign s^* to one of these simplices, and the relevant incidences will then be counted in the subproblem associated with that simplex; see below for details.

(ii) Consider next incidences involving dual points that lie on some 2-face f of some simplices of Ξ. In primal 4-space, the 2-flat π spanned by f is mapped to a line ℓ, such that any dual hyperplane p^* that fully contains f (that is, π) is mapped back to a point in primal 4-space that lies in ℓ. Intersecting ℓ with the paraboloid $w = x^2 + y^2 + z^2$ and projecting down to the original 3-space, we get at most two original

points p whose dual images p^* can fully contain f. Hence, the number of incidences that fall into this special case are at most $\sum_f 2|S^* \cap f|$, over all 2-faces f as above.

(iii) For f, π, and ℓ as in (ii), the dual hyperplanes p^* that do not fully contain f must cross every simplex that has f as a face. Indeed, let σ be such a simplex. Since p^* meets f, it intersects the closure $\bar{\sigma}$ of σ. Since p^* does not cross σ, it must be a supporting hyperplane to $\bar{\sigma}$. But such a supporting hyperplane must meet $\bar{\sigma}$ in a full face of some dimension. Hence, its intersection with f must be at some subface of f, contrary to assumption. Therefore, as in step (i), we assign each dual point $s^* \in f$ to one of these adjacent simplices, and the relevant incidences will then be counted in the subproblem associated with that simplex.

Getting back to the bound $\sum_f 2|S^* \cap f|$ in (ii) for the incidences with hyperplanes that fully contain f. Dual points s^* that lie on exactly one 2-face f contribute a total of $O(n)$ to the sum. Consider then a dual point s^* that lies on two (or more) 2-faces f_1, f_2. If f_1, f_2 are co-planar, we can ignore one of them, because the two hyperplanes containing f_1 and the two containing f_2 are the same (see the argument in step (ii) above, which only depends on the 2-flat supporting f_1 (which is the same as the one supporting f_2), and not on f_1 (or f_2) itself), so we still have at most two incidences involving s^*. If f_1 and f_2 are not coplanar, then one of the hyperplanes containing f_2 must cross f_1, so the two incidences of s^* within f_2 can be charged to this crossing incidence, which is handled as above, by assigning s^* into one of the simplices bounded by f_2. It is easily checked that the same argument applies when s^* lies on more than two 2-faces: every additional 2-face (which is not coplanar with f_1) will be contained in a dual hyperplane that crosses f_1 (or else not contribute any new incidence), so the corresponding incidences can be charged to a suitable crossing incidence, as above. To summarize, the sum $\sum_f 2|S^* \cap f|$ is $O(n)$ plus an excess that will be handled within the simplices of Ξ.

(iv) Consider next incidences involving dual points on some edge e of some simplices of Ξ. In primal 4-space, the line ℓ spanned by e is mapped to a 2-flat π, such that any dual hyperplane p^* that fully contains e (that is, ℓ) is mapped back to a point in primal 4-space that lies in π. Intersecting π with the paraboloid $w = x^2 + y^2 + z^2$ and projecting down to the original 3-space, we conclude that the original point p that is mapped to p^*, for any p^* as above, lies in the intersection of V with a circle c. Similarly, any dual point s^* that lies in e (and thus in ℓ) is mapped in primal 4-space to a hyperplane that contains π. Intersecting that hyperplane with the paraboloid, and projecting down to the original 3-space, we obtain a sphere s that fully contains c. We may assume that c is not fully contained in V, because otherwise, incidences between points on c and spheres that contain c are already recorded in the complete bipartite graph $P_c \times S_c$ that we have removed from $G(P, S)$. But then $|P \cap c| \leq 2D = O(1)$. That is, at most $2D$ dual hyperplanes fully contain e, yielding at most $2D|S^* \cap e|$ incidences, for a total, over all edges e, of $\sum_e 2D|S^* \cap e|$ incidences.

(v) Arguing as in step (iii), dual hyperplanes p^* that cross e but do not fully contain it must cross every simplex that has e as an edge. Hence, an incidence of such a dual hyperplane p^* with a dual point on e can be charged to a crossing of some adjacent simplex by p^*, and any such hyperplane-simplex crossing can be charged only $O(1)$ times—at most once for each edge of the simplex being crossed by p^*. In total, we get a total of $O(r^4 \cdot (m/r)) = O(mr^3)$ such incidences.

Again, in the sum $\sum_e 2D|S^* \cap e|$, obtained in step (iv), dual points s^* that lie on just one simplex edge e contribute at most $2Dn = O(n)$ incidences. Consider then a point s^* that lies on more than one edge, say on edges e_1 and e_2. In the primal space, they correspond to distinct circles c_1, c_2, both contained in s, and any point p on $c_2 \setminus c_1$ corresponds to a dual hyperplane that (contains e_2 and) crosses e_1, so the incidences with the dual hyperplanes that contain e_2 can be charged to crossing incidences involving e_1, as above. This also works for any number of edges containing s^*. Hence, $\sum_e 2D|S^* \cap e|$ is $O(n)$ plus a term proportional to the number of "crossing incidences", which, as has just been argued, is $O(mr^3)$.

(vi) Consider finally incidences that involve dual points that are vertices of some simplices of Ξ. If such a vertex s^* does not lie in the relative interior of any higher-dimensional face of any other simplex, that is, all the simplices adjacent to s^* have s^* as a vertex, then any incidence between a dual hyperplane p^* and

s^* can be charged to the crossing of p^* with some simplex σ of Ξ that is adjacent to s^*. It follows, arguing as in (v), that the number of incidences of this kind is at most $O(r^4) \cdot m/r = O(mr^3)$. On the other hand, if s^* lies in the relative interior of some higher-dimensional face f of some other simplex, we handle the incidence between s^* and any hyperplane p^* as in steps (i)–(v) above.

To recap, ignoring incidences that have been assigned to the subproblems within the simplices of Ξ, as well as incidences that have been recorded in the complete bipartite graphs $P_c \times S_c$, we have accumulated in this step only $O(n + mr^3)$ incidences.

Incidences within the simplices of Ξ. We now proceed to consider incidences within the simplices of Ξ. For each simplex σ of Ξ, let n_σ denote the number of points of S^* in the interior of σ, including the points that have been assigned to σ from its boundary as above. We bound, for each simplex σ of Ξ, the number of incidences between the n_σ dual points in its interior and the at most m/r dual hyperplanes that cross σ. By duplicating simplices σ for which $n_\sigma > n/r^4$, so that in each copy we take at most n/r^4 of these points (but retain all crossing hyperplanes), we obtain a collection of $O(r^4)$ simplices, each of which is crossed by at most m/r dual hyperplanes and contains at most n/r^4 dual points; we denote the actual number of these hyperplanes and points as m_σ and n_σ (the latter notation is slightly abused, as it now refers only to a single copy (subproblem) of σ), respectively, for each simplex σ.

For each cell σ, we apply the bound (5) to the subset $P^{(\sigma)}$ of the points of P whose dual hyperplanes cross σ, and to the subset $S^{(\sigma)}$ of the spheres whose dual points lie in σ, and note that the case $m/r = O((n/r^4)^{1/4})$ does not arise, because then we would also have $m = O(n^{1/4})$, and then we would have used instead the bound (6), avoiding the partitioning altogether. That is, we get, for each σ,

$$I(P^{(\sigma)}, S^{(\sigma)}) = O\left(m_\sigma^{2/3} n_\sigma^{2/3} + m_\sigma + n_\sigma^{3/2} \log^\kappa n_\sigma + \sum_c |P_c^{(\sigma)}| \cdot |S_c^{(\sigma)}|\right),$$

for a suitable complete bipartite decomposition $\bigcup_c \left(P_c^{(\sigma)} \times S_c^{(\sigma)}\right)$.

We sum these bounds, over the simplices σ of Ξ. We note that the same circle c may arise in many complete bipartite graphs $P_c^{(\sigma)} \times S_c^{(\sigma)}$, but (i) all these graphs are contained in $P_c \times S_c$, and (ii) they are edge disjoint, because each dual point s^* lies in (or is a boundary point which is assigned to) at most one simplex. This allows us to replace all the partial subgraphs $P_c^{(\sigma)} \times S_c^{(\sigma)}$ by the single graph $P_c \times S_c$, for each circle c (contained in V). We thus get

$$\begin{aligned}
I(P,S) &= O\left(\sum_\sigma \left(m_\sigma^{2/3} n_\sigma^{2/3} + m_\sigma + n_\sigma^{3/2} \log^\kappa n_\sigma\right) + mr^3 + n + \sum_c |P_c| \cdot |S_c|\right) \\
&= O\left(r^4\left((m/r)^{2/3}(n/r^4)^{2/3} + (n/r^4)^{3/2} \log^\kappa (n/r^4)\right) + mr^3 + n + \sum_c |P_c| \cdot |S_c|\right) \\
&= O\left(m^{2/3} n^{2/3} r^{2/3} + \frac{n^{3/2}}{r^2} \log^\kappa (n/r^4) + mr^3 + n + \sum_c |P_c| \cdot |S_c|\right). \qquad (7)
\end{aligned}$$

We now choose $r = \dfrac{n^{5/16} \log^{3\kappa/8}(m^4/n)}{m^{1/4}}$, to equalize (asymptotically) the first two terms in the bound (7), which then become $O(m^{1/2} n^{7/8} \log^{\kappa/4}(m^4/n))$. The third term becomes $mr^3 = m^{1/4} n^{15/16} \log^{9\kappa/8}(m^4/n)$, which is dominated by the preceding bound for $m = \Omega(n^{1/4})$, as is easily checked. As already noted, the complementary case $m = O(n^{1/4})$ has been handled by (6), and the case $m = \Omega(n^{5/4} \log^{3\kappa/2} n)$ is handled simply by (5) (now without the term $n^{3/2} \log^\kappa n$ as it is subsumed by the other term).

This completes the proof of Theorem 1.1. □

Remark. In retrospect, once we have reduced the problem to that of bounding the number $I(P^*, \Gamma^*)$ on incidences between the projected points and curves on the xy-plane, we could have applied, as a black-box, the analysis of Sharir and Zahl [38], and get a slightly weaker bound with an additional (arbitrarily small) ε in the exponents.

As this remark will be significant in the proofs of some of the forthcoming applications, we rephrase here the result of [38], with the notation used in the proof of Theorem 1.1, for the convenience of the reader.

Theorem 2.1 (Sharir and Zahl [38])**.** *Let Γ^* be a set of n algebraic plane curves that belong to an s-dimensional family F of curves of maximum constant degree D, no two of which share a common irreducible component, and let P^* be a set of m points in the plane. Then, for any $\varepsilon > 0$, the number $I(P^*, \Gamma^*)$ of incidences between the points of P^* and the curves of Γ^* satisfies*

$$I(P^*, \Gamma^*) = O\Big(m^{\frac{2s}{5s-4}} n^{\frac{5s-6}{5s-4}+\varepsilon} + m^{2/3}n^{2/3} + m + n\Big),$$

where the constant of proportionality depends on ε, s, D, and the complexity of the family F.

In this result, an *s-dimensional family of curves* is a family $\mathcal{C}$ of algebraic curves (of constant maximum degree), so that each curve in $\mathcal{C}$ can be represented as a point in some finite-dimensional parametric space, and the set of these "dual" points is an s-dimensional algebraic variety of constant degree (which is referred to as the "complexity" of F).

In our case, $s = 4$, since each curve of Γ^* can be represented by the four parameters that define the corresponding sphere, and, using the fact that V is of constant degree, it is easy to verify that the assumptions of Theorem 2.1 hold in this case. Substituting $s = 4$ gives us our bound, except that the polylogarithmic factor is replaced by the factor n^{ε}. That is, exploiting the fact that we are dealing here with spheres, so that the dual representation involves points and *hyperplanes*, allows us to obtain the finer bound in (2) (concretely, using cuttings instead of the rather involved partitioning scheme of [31]).

3 Distinct and repeated distances in three dimensions

In this section we prove Theorems 1.2 and 1.3, the applications of our main result to distinct and repeated distances in three dimensions. The theorems present four results, in each of which the problem is reduced to one involving incidences between spheres and points on a surface. However, except for Theorem 1.2(b), the spheres that arise in the other cases are restricted, by requiring their centers to lie on V and / or to have a fixed radius. This makes the spheres have only three or two degrees of freedom. The case of two degrees of freedom (in Theorem 1.3(a)) is the simplest, and requires very little of the machinery developed here (see below). The cases of three degrees of freedom (in Theorem 1.2(a) and Theorem 1.3(b)) call for a dual representation of the setup in three dimensions.

A rigorous analysis along this line is doable, and we will comment on it later, but there are several technical issues that arise, and a careful treatment of them will make the proofs longer and somewhat more involved. As a compromise, we state the sharp bounds that would result from the full analysis, but present simpler "black-box" proofs that exploit the machinery in [38] and yield slightly inferior bounds.

Proof of Theorem 1.2. We will first establish the more general bound in (b); handling (a) requires a somewhat different analysis.

(b) Let t denote the number of distinct distances in $P_1 \times P_2$. For each $q \in P_2$, draw t spheres centered at q and having as radii the t distinct distances. We get a collection S of nt spheres, a set P_1 of m points

on V, which we relabel as P, to simplify the notation, and exactly mn incidences between the points of P and the spheres of S.

In order to effectively apply the bound in Theorem 1.1, we first have to control the term $\sum_i |P_i| \cdot |S_i|$; that is, we argue that most of the mn incidences do not come from this bound, unless $t = \Omega(n)$. Indeed, for each i, we have $|S_i| \leq 2t$; this is because all the spheres in S_i pass through a fixed circle c, so, up to multiplicity 2, their radii are all distinct. This implies that

$$\sum_i |P_i| \cdot |S_i| \leq 2t \sum_i |P_i| = O(mt).$$

If this would have accounted for more than, say, half the incidences, we would get $t = \Omega(n)$, as claimed, and then the bound in the theorem would follow. We may thus ignore this term, and write

$$mn = O\left(m^{2/3}(nt)^{2/3} + m^{1/2}(nt)^{7/8}\log^{\beta}(m^4/n) + m + nt\right),$$

or

$$t = \Omega\left(\min\left\{m^{1/2}n^{1/2},\ m^{4/7}n^{1/7}/\log^{8\beta/7}(m^4/n),\ m\right\}\right),$$

as claimed.

(a) Here we are in a more favorable situation, because the spheres in S have only three degrees of freedom, in the sense that their centers lie on the surface V, so that, in principle, we need only two parameters to specify the center and one for the radius.

One possibility would be to adapt the analysis in the proof of Theorem 1.1, with the difference that the spheres are now dualized to points in three dimensions, rather than four. As already noted, this would raise several technical issues, which, albeit minor, require careful analysis that would be too space-consuming. A discussion of the issues that arise and the way to handle them to get the sharper bound is given below.

Instead, we "shortcut" the analysis, and apply the improved incidence bound of Sharir and Zahl [38], stated in Theorem 2.1, with $s = 3$. That is, we still represent each curve γ_s^* of Γ^* by the parameters $(x, y, z, r) \in \mathbb{R}^4$ that define the corresponding sphere s (where (x, y, z) is its center and r its radius), but now (x, y, z) is constrained to lie on V. It then easily follows that Γ^* is a three-dimensional family of curves (in the notation of Theorem 2.1).

We thus get the bound

$$I(P^*, \Gamma^*) = O\left(|P^*|^{2/3}|\Gamma^*|^{2/3} + |P^*|^{6/11}|\Gamma^*|^{9/11+\varepsilon} + |P^*| + |\Gamma^*|\right),$$

for any $\varepsilon > 0$. Arguing as in the proof of (b), we may ignore the term $\sum_i |P_i| \cdot |S_i|$ in the bound on $I(P, S)$, which is negligible unless $t = \Omega(n)$, and thus get the inequality

$$n^2 = O\left(n^{2/3}(nt)^{2/3} + n^{6/11}(nt)^{9/11+\varepsilon} + nt\right),$$

which yields $t = \Omega(n^{7/9-\varepsilon})$, for any $\varepsilon > 0$, thereby completing the proof of (the coarser version of) (a). □

Proof of Theorem 1.3.

Consider (a) first. Following the standard approach to problems involving repeated distances, we draw a unit sphere s_p around each point $p \in P$, and seek an upper bound on the number of incidences between these spheres and the points of P; this latter number is exactly twice the number of unit distances determined by P.

This instance of the problem has several major advantages over the general analysis in Theorem 1.1. First, in this case the incidence graph $G(P, S)$ cannot contain $K_{3,3}$ as a subgraph, eliminating altogether

the complete bipartite graph decomposition in (1) (or, rather, bounding the overall number of edges in these subgraphs by $O(n)$).

More importantly, the family Γ^* of curves "almost" has only two degrees of freedom. To have two degrees of freedom, in the sense of Pach and Sharir [32], it is required that, for any pair of points $p^*, q^* \in P^*$, there are at most $O(1)$ curves of Γ^* passing through p^* and q^* (and that any pair of curves of Γ^* intersect in at most $O(1)$ points, a property that we have already established).

To test for this property, fix a pair $p^*, q^* \in P^*$. By our assumption that the coordinate frame is generic, there is a unique pair $p, q \in P$ that project, respectively, to p^* and q^*, and any curve $\gamma_s^* \in \Gamma^*$ that passes through p^* and q^* is the projection of a unique curve $\gamma_s \in \Gamma$ that passes through p and q. The corresponding sphere $s \in S$ is then a unit sphere that passes through p and q, so its center must lie on a suitable circle c_{pq} that is centered at $\frac{1}{2}(p+q)$ and is orthogonal to $\vec{pq}$. As is easily checked, the circles c_{pq} are all distinct.

If c_{pq} is not fully contained in V, it meets it in at most $2D$ points, implying that there are at most $2D = O(1)$ curves of Γ^* that pass through p^* and q^*, as desired.

It remains to study pairs p^*, q^* for which c_{pq} is fully contained in V. The number of curves that pass through p^* and q^* is $|P \cap c_{pq}|$. By the result of Sharir, Sheffer, and Zahl [35], already mentioned in the proof of Theorem 1.1, except for two popular points (which we may assume, as above, not to belong to P), every point $p \in P$ is incident to at most $44D^2 = O(1)$ circles that are fully contained in V. It follows that

$$\sum_{p,q \in P} |P \cap c_{pq}| = O(n).$$

We can now apply Székely's crossing lemma argument [40] to P^* and Γ^*. The edges in Székely's graph have constant multiplicity, except for those that connect pairs p^*, q^* for which $c_{pq} \subset V$. As just argued, the overall number of edges of the latter kind is $O(n)$. Omitting these edges from the graph, Székely's argument applies to the remainder, and yields the bound $O(n^{4/3})$ for the number of edges. Combining this bound with the linear bound on the number of high-multiplicity edges, and the additional linear bound on the size of the complete bipartite graphs $P_i \times S_i$, as noted above, we get a total of $O(n^{4/3})$ incidences, and thus $O(n^{4/3})$ unit distances.

We now consider (b). Again, we reduce the problem to that of bounding the number of incidences between the m points of P_1, which lie on V, and the n unit spheres centered at the points of P_2. Here too the overall number of edges in the complete bipartite graph decomposition is $O(m+n)$, so we can ignore this part of the bound.

In this case, the curves of Γ^* have three degrees of freedom, or, in the terminology of Sharir and Zahl [38], as reviewed in Theorem 2.1, Γ is a three-dimensional family of curves. Applying the same reasoning as in this preceding proof, we conclude that the number of unit distances in this case is

$$O\left(m^{2/3}n^{2/3} + m^{6/11}n^{9/11+\varepsilon} + m + n\right),$$

for any $\varepsilon > 0$. □

Improving the bounds. In the proofs of Theorem 1.2(a) and Theorem 1.3(b), we want to dualize the problem in a way that exploits the fact that the spheres of S have only three degrees of freedom. We still map the spheres to points in $\mathbb{R}^4$ and the points to hyperplanes in $\mathbb{R}^4$, as above, but now the dual points s^* all lie on a three-dimensional algebraic variety $V^* \subset \mathbb{R}^4$ of constant degree; in Theorem 1.2(a), $V^* = V \times \mathbb{R}$, and in Theorem 1.3(b), V^* is the paraboloid $x_4 = x_1^2 + x_2^2 + x_3^2 + 1$. We construct a $(1/r)$-cutting for the collection of the dual hyperplanes, but use a coarser (and simpler) technique of drawing a random sample of $O(r \log r)$ hyperplanes, construct their arrangement, and triangulate each cell into simplices. We now use

the generalized zone theorem of Aronov et al. [7], to conclude that V^* crosses only $O(r^3 \log^4 r)$ simplices, and we apply the weak bound only in these simplices. There are various additional technical issues that have to be handled, but, as already explained above, we skip over them, in the interest of keeping the paper short. Working out all the details, we get the slightly improved bounds, as asserted in the theorems.

4 Proof of Theorem 1.4

The proof establishes the bound in (4), via induction, adding a prespecified approximation parameter $\varepsilon > 0$ to the bound. Concretely, we claim that, for any prespecified $\varepsilon > 0$ we can write

$$G(P,S) = G_0(P,S) \cup \bigcup_c (P_c \times S_c),$$

so that

$$|G_0(P,S)| \leq A\left(m^{8/11+\varepsilon} n^{9/11} + m + n\right), \tag{8}$$

and

$$\sum_c \left(|P_c| + |S_c|\right) = O\left(m^{8/11+\varepsilon} n^{9/11} + m + n\right),$$

where A and the other constants of proportionality depend on ε.

The base cases are when $m \leq n^{1/4}$ or when $m \leq m_0$, for some sufficiently large constant m_0 that will be set later.

Consider first the case $m \leq n^{1/4}$. We adapt the argument for this case given in the proof of Theorem 1.1. It yields the bound $I(P, S_1) = O(n)$ for the set S_1 of spheres of S that are not strongly degenerate. Each strongly degenerate sphere s can be replaced by the unique circle c_s that contains all its incident points. We thus get an incidence problem involving a set P of m points and a multiset $\mathcal{C}$ of n circles in $\mathbb{R}^3$.

Fix the threshold multiplicity $\mu_0 = n^{1/4}$, and consider the set $\mathcal{C}^-$ of all circles in $\mathcal{C}$ with multiplicity at most μ_0. Each circle $c \in \mathcal{C}^-$ with at most two points of P on it contributes at most $2\mu(c)$ incidences, where $\mu(c)$ denotes the multiplicity of c. Summing these bounds over all such circles, we get at most $2|S| = 2n$ incidences. The number of incidences involving circles containing at least three points of P is $O(m^3 \cdot \mu_0) = O(n)$.

This leaves us with circles of multiplicity larger than $n^{1/4}$. We represent the corresponding incident pairs as a union of complete bipartite graphs $P_c \times S_c$, over all circles in $\mathcal{C}$ with multiplicity larger than $n^{1/4}$. We clearly have $\sum_c |S_c| \leq n$, and $\sum_c |P_c|$ is simply the number of incidences between the points of P and the at most $n^{3/4}$ "heavy" circles, counted *without* multiplicity. The same argument used above gives the bound $O(m^3 + n^{3/4}) = O(n^{3/4})$.

In summary, we have for this case

$$I(P,S) = O\left(n + \sum_c |P_c| \cdot |S_c|\right), \tag{9}$$

where $\bigcup_c (P_c \times S_c)$ is contained in the incidence graph $G(P,S)$, and $\sum_c |P_c| = O(n^{3/4})$ and $\sum_c |S_c| = O(n)$. That is, (8) holds in this case.

The case $m \leq m_0$ follows easily if we choose A sufficiently large. This holds for any choice of m_0; the value that we choose is specified later.

Suppose then that (8) holds for all sets P', S', with $|P'| < m$, $|S'| < n$, and consider the case where the sets P, S are of respective sizes m, n, and $m > n^{1/4}$ and $m > m_0$.

Before continuing, we also dispose of the case $m \geq n^3$. In this case we consider the arrangement $\mathcal{A}(S)$ of the spheres in S. The complexity of $\mathcal{A}(S)$ is $O(n^3)$. More precisely, this bound holds, and is asymptotically tight, for spheres in general position. In our case, S is likely not to be in general position, and then the complexity of $\mathcal{A}(S)$ might be smaller, because vertices and edges might be incident to many spheres. Nevertheless, if we count each vertex and edge of $\mathcal{A}(S)$ with its multiplicity, we still get the upper complexity bound $O(n^3)$.

This means that the number of incidences with points that are either vertices or lie on the (relatively open) faces of $\mathcal{A}(S)$ is $O(n^3) = O(m)$. Incidences with points that lie on the (relatively open) edges of $\mathcal{A}(S)$ (note that each such edge is a portion of some circle) are recorded, as usual, by a complete bipartite graph decomposition $\bigcup_c (P_c \times S_c)$, where, as just argued, we have $\sum_c |P_c| \leq m$ and $\sum_c |S_c| = O(n^3) = O(m)$. This implies that (8) holds in this case, so, in what follows, we assume that $m \leq n^3$.

Applying the polynomial partitioning technique. We fix a sufficiently large constant parameter $D \ll m^{1/3}$, whose choice will be specified later, and apply the polynomial partitioning technique of Guth and Katz [24]. We obtain a polynomial $f \in \mathbb{R}[x, y, z]$ of degree at most D, whose zero set $Z(f)$ partitions 3-space into $O(D^3)$ (open) connected components (cells), and each cell contains at most $O(m/D^3)$ points. By duplicating cells if necessary, we may also assume that each cell is crossed by at most $O(n/D)$ spheres of S; this duplication keeps the number of cells $O(D^3)$ (because each sphere crosses only $O(D^2)$ cells). That is, we obtain at most aD^3 subproblems, for some absolute constant a, each associated with some cell of the partition, so that, for each $i \leq aD^3$, the i-th subproblem involves a subset $P_i \subset P$ and a subset $S_i \subset S$, such that $m_i := |P_i| \leq bm/D^3$ and $n_i := |S_i| \leq bn/D$, for another absolute constant b.

Set $P_0 := P \cap Z(f)$ and $P' = P \setminus P_0$. We have

$$I(P, S) = I(P_0, S) + I(P', S). \tag{10}$$

We first bound $I(P_0, S)$. Decompose $Z(f)$ into its $O(D)$ irreducible components, assign each point of P_0 to every component that contains it, and assign the spheres of S to all components. We now fix a component, and bound the number of incidences between the points and spheres assigned to that component; $I(P_0, S)$ is at most D times the bound that we get.

We may therefore assume that $Z(f)$ is irreducible. If $Z(f)$ is a plane or a sphere, then for any sphere $s \in S$, the curve $s \cap Z(f)$ is a circle; let $\mathcal{C}$ denote the multiset of these circles, where each circle has multiplicity equal to the number of spheres that contain it. Then $I(P_0, S)$ is the number of incidences between the points of P_0 and the circles of $\mathcal{C}$, counted with multiplicity.

We bound the number of incidences of this latter kind using the incidence bound of Aronov et al. [9] for points and circles in $\mathbb{R}^3$. Fixing a threshold μ, the number of incidences involving circles with multiplicity at most μ (and counted with their multiplicity) is easily seen to be

$$O\left(m^{2/3}n^{2/3}\mu^{1/3} + m^{6/11}n^{9/11}\mu^{2/11}\log^{2/11}(m^3\mu/n) + m\mu + n\right). \tag{11}$$

We now choose

$$\mu = \min\left\{m^{2/11}n^{5/11},\ m,\ n^{9/11}/m^{3/11}\right\}.$$

An easy, albeit a bit tedious, calculation shows that the bound in (11) becomes $O(m^{8/11}n^{9/11} + n)$.

For circles c with multiplicity larger than μ, we record the corresponding point-sphere incident pairs by a complete bipartite graph decomposition $\bigcup_c (P_c \times S_c)$, where c ranges over all such "heavy" circles, and where $P_c = P_0 \cap c$ and S_c is the set of all spheres that contain c. We clearly have $\sum_c |S_c| = O(n)$ (each sphere can intersect $Z(f)$ in only one circle, except for the unique sphere, if any, that coincides with $Z(f)$, which we may ignore), and $\sum_c |P_c|$ is the number of incidences between the points of P_0 and the heavy

circles, counted *without* multiplicity. The number of these circles is at most $O(n/\mu)$. Using the bound in [9], we get, as above,

$$\sum_c |P_c| = O\left(\frac{m^{2/3}n^{2/3}}{\mu^{2/3}} + \frac{m^{6/11}n^{9/11}}{\mu^{9/11}}\log^{2/11}(m^3\mu/n) + m + \frac{n}{\mu}\right).$$

Since this is asymptotically the same as the bound in (11) divided by μ, we simply (and pessimistically) upper bound this by $O(m^{8/11}n^{9/11} + n)$.

Assume then that $Z(f)$ is neither a plane nor a sphere. Since $\deg(Z(f)) \leq D$ is a constant, our main Theorem 1.1 implies that

$$I(P_0, S) = O\left(m^{2/3}n^{2/3} + m^{1/2}n^{7/8}\log^{\beta}(m^4/n) + m + n + \sum_c |P_c| \cdot |S_c|\right),$$

where $\bigcup_c (P_c \times S_c) \subseteq G(P_0, S)$, and $\sum_c |P_c| = O(m)$ and $\sum_c |S_c| = O(n)$. As is easily checked, the first four terms are subsumed in (8), if we choose A sufficiently large, and the term $\sum_c |P_c| \cdot |S_c|$ is added to the complete bipartite graph decomposition that we collect.

Finally, we estimate

$$I(P', S) = \sum_{i=1}^{aD^3} I(P_i, S_i).$$

By the induction hypothesis, we get

$$I(P_i, S_i) \leq A\left(m_i^{8/11+\varepsilon}n_i^{9/11} + m_i + n_i + \sum_c |P_{i,c}| \cdot |S_{i,c}|\right),$$

for a suitable complete bipartite decomposition $\bigcup_c (P_{i,c} \times S_{i,c}) \subseteq G(P_i, S_i)$.

When summing these bounds, we note that the same circle c may arise in many complete bipartite graphs, but, as already noted earlier, (i) all these graphs are contained in $P_c \times S_c$, and (ii) they are edge disjoint, because each point $p \in P'$ lies in at most one cell, and even if this cell gets duplicated, the relavant spheres are all distinct. This allows us to replace all the partial subgraphs $P_{i,c} \times S_{i,c}$ by the single graph $P_c \times S_c$, for each circle c in the decomposition.

The sum of the other terms is

$$\begin{aligned} I_0(P', S) &\leq A \cdot aD^3\left((bm/D^3)^{8/11+\varepsilon}(bn/D)^{9/11} + (bm/D^3) + (bn/D)\right) \\ &= \frac{Aab^{17/11+\varepsilon}}{D^{3\varepsilon}}m^{8/11+\varepsilon}n^{9/11} + Aabm + AabD^2n. \end{aligned}$$

We note that $m^{8/11+\varepsilon}n^{9/11} \geq m^{\varepsilon} \cdot m$ and $m^{8/11+\varepsilon}n^{9/11} \geq m^{\varepsilon} \cdot n$ for for $n^{1/4} \leq m \leq n^3$. We choose D sufficiently large so that $D^{3\varepsilon} \geq 4ab^{17/11+\varepsilon}$, and then the bound is at most

$$\left(\frac{A}{4} + \frac{Aab}{m^{\varepsilon}} + \frac{AabD^2}{m^{\varepsilon}}\right)m^{8/11+\varepsilon}n^{9/11}.$$

Choosing m_0 sufficiently large, so that $m_0^{\varepsilon} \geq 4abD^2$, we ensure that, for $m \geq m_0$,

$$\frac{A}{4} + \frac{Aab}{m^{\varepsilon}} + \frac{AabD^2}{m^{\varepsilon}} \leq \frac{A}{4} + \frac{A}{4} + \frac{A}{4} = \frac{3A}{4}.$$

Adding the bounds for $I(P_0, S)$, and choosing A sufficiently large, we get that (8) hold for P and S. The coorresponding bounds on $\sum_c |P_c|$ and $\sum_c |S_c|$ are established by the same inductive analysis, and we omit the straightforward details. This establishes the induction step and thereby completes the proof. □

5 Discussion

In this paper we have made significant progress on a major incidence problem involving points and spheres in three dimensions, for the special case where the points lie on a constant-degree algebraic surface. We have also obtained several applications of this result to problems involving repeated and distinct distances in three dimensions, and have extended the analysis to the case where the points are arbitrary and are not required to lie on a constant-degree surface; this latter extension improves the bound derived in Apfelbaum and Sharir [5], and it is also significantly more general, as it does not require the spheres to be non-degenerate, as in [5].

The study in this paper raises several interesting open problems.

(i) Our analysis suggests that if, instead of the family of spheres, we take S to be a *k-dimensional family* of constant-degree algebraic surfaces (in the terminology of [38], already mentioned above), and still assume the points to lie on a constant-degree surface, we can then extend the analysis in Theorem 1.1 to get an analogous bound for point-surface incidences, depending on k, and resembling the one obtained in [38] for point-curve incidences in the plane.

It also seems likely that, as in Theorem 1.4, the analysis can be further extended to the case where the points do not have to lie on a surface, and that the bound that it yields is

$$O\left(m^{\frac{2s}{3s-1}+\varepsilon}n^{\frac{3s-3}{3s-1}} + m + n + \sum_{\gamma} |P_\gamma| \cdot |S_\gamma|\right),$$

for any $\varepsilon > 0$, where $\bigcup_\gamma (P_\gamma \times S_\gamma) \subseteq G(P,S)$, and where $\sum_\gamma |P_\gamma|$, $\sum_\gamma |S_\gamma|$ are suitably bounded. Assuming that this extension can be made rigorous, it would yield a significant generalization of Zahl's result [42], where the leading term almost matches his bound, but there are no restrictions that the incidence graph does not contain a fixed-size complete bipartite subgraph, as assumed in [42].

(ii) A long-standing open problem is that of establishing the lower bound of $\Omega(n^{2/3})$ for the number of distinct distances determined by a set of n points in $\mathbb{R}^3$, without assuming them to lie on a surface. The best known lower bound is due to Solymosi and Vu [39]. In the present study we have obtained some partial results (with better lower bounds) for cases where the points do lie on a surface. We hope that some of the ideas used in this work could be applied for the general problem.

(iii) Another major long-standing open problem is that of improving the upper bound $O(n^{3/2})$, established in [29, 42], on the number of unit distances determined by a set of n points in $\mathbb{R}^3$, again without assuming them to lie on a surface. It would be interesting to make progress on this problem.

(iv) Finally, it would be interesting to find additional applications of the results of this paper. One direction to look at is the analysis of repeated patterns in a point set, such as congruent or similar simplices, which can sometimes be reduced to point-sphere incidence problems; see [2, 3].

References

[1] P. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir and S. Smorodinsky, Lenses in arrangements of pseudocircles and their applications, *J. ACM* 51 (2004), 139–186.

[2] P. K. Agarwal, R. Apfelbaum, G. Purdy and M. Sharir, Similar simplices in a d-dimensional point set, *Proc. 23rd Annu. ACM Sympos. Comput. Geom.* (2007), 232–238.

[3] P. Agarwal and M. Sharir, On the number of congruent simplices in a point set, *Discrete Comput. Geom.* 28 (2002), 123–150.

[4] R. Apfelbaum and M. Sharir, Large bipartite graphs in incidence graphs of points and hyperplanes, *SIAM J. Discrete Math.* 21 (2007), 707–725.

[5] R. Apfelbaum and M. Sharir, Non-degenerate spheres in three dimensions, *Combinat. Probab. Comput.* 20 (2011), 503–512.

[6] B. Aronov, J. Pach, M. Sharir and G. Tardos, Distinct distances in three and higher dimensions, *Combinat. Probab. Comput.* 13 (2004), 283–293.

[7] B. Aronov, M. Pellegrini and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9 (1993), 177–186.

[8] B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds for incidences, *Discrete Comput. Geom.* 28 (2002), 475–490.

[9] , B. Aronov, V. Koltun and Micha Sharir, Incidences between points and circles in three and higher dimensions, *Discrete Comput. Geom.* 33.2 (2005), 185–206.

[10] A. Basit and A. Sheffer, Incidences with k-non-degenerate sets and their applications, *J. Comput. Geom.* 5 (2014), 284–302.

[11] S. Basu and M. Sombra, Polynomial partitioning on varieties of codimension two and point-hypersurface incidences in four dimensions, *Discrete Comput. Geom.* 55.1 (2016), 158–184.

[12] P. Brass, Exact point pattern matching and the number of congruent triangles in a threedimensional point set, *Proc. European Sympos. Algorithms*, 2000, Springer LNCS 1879, pp. 112–119.

[13] P. Brass and Ch. Knauer, On counting point-hyperplane incidences, *Comput. Geom. Theory Appls.* 25 (2003), 13–20.

[14] P. Brass, W. Moser and J. Pach, *Research Problems in Discrete Geometry*, Springer Verlag, New York, 2005.

[15] B. Chazelle, Cuttings, in *Handbook of Data Structures and Applications* (D. P. Mehta and S. Sahni, Eds.), Chapman and Hall/CRC, 2004.

[16] F. R. K. Chung, Sphere-and-point incidence relations in high dimensions with applications to unit distances and furthest-neighbor pairs, *Discrete Comput. Geom.* 4 (1989), 183–190.

[17] J. Cilleruelo, A. Iosevich, B. Lund, O. Roche-Newton and M. Rudnev, Elementary methods for incidence problems in finite fields, in arXiv:1407.2397.

[18] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.

[19] H. Edelsbrunner, L. Guibas and M. Sharir, The complexity of many cells in arrangements of planes and related problems, *Discrete Comput. Geom.* 5 (1990), 197–216.

[20] J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl, A semi-algebraic version of Zarankiewicz's problem, *J. European Math. Soc.*, to appear. Also in arXiv:1407.5705.

[21] W. Fulton, *Introduction to Intersection Theory in Algebraic Geometry*, Expository Lectures from the CBMS Regional Conference Held at George Mason University, June 27–July 1, 1983, Vol. 54. AMS Bookstore, 1984.

[22] L. Guth, Polynomial partitioning for a set of varieties, in arXiv:1410.8871.

[23] L. Guth and N. H. Katz, Algebraic methods in discrete analogs of the Kakeya problem, *Advances Math.* 225 (2010), 2828–2839. Also in arXiv:0812.1043v1.

[24] L. Guth and N. H. Katz, On the Erdős distinct distances problem in the plane, *Annals Math.* 181 (2015), 155–190. Also in arXiv:1011.4105.

[25] L. Guth and J. Zahl, Algebraic curves, rich points, and doubly-ruled surfaces, in arXiv:1503.02173.

[26] C. G. A. Harnack, Über die Vieltheiligkeit der ebenen algebraischen Curven, *Math. Ann.* 10 (1876), 189–199.

[27] J. Harris, *Algebraic Geometry: A First Course*, Vol. 133, Springer-Verlag, New York, 1992.

[28] H. Kaplan, J. Matoušek and M. Sharir, Simple proofs of classical theorems in discrete geometry via the Guth–Katz polynomial partitioning technique, *Discrete Comput. Geom.*, submitted. Also in arXiv:1102.5391.

[29] H. Kaplan, J. Matoušek, Z. Safernová and M. Sharir, Unit distances in three dimensions, *Combinat. Probab. Comput.* 21 (2012), 597–610. Also in arXiv:1107.1077.

[30] A. Marcus and G. Tardos, Intersection reverse sequences and geometric applications, *J. Combinat. Theory* Ser. A 113 (2006), 675–691.

[31] J. Matoušek and Z. Patáková, Multilevel polynomial partitioning and simplified range searching, in arXiv:1406.3058v2 (July 2015).

[32] J. Pach and M. Sharir, On the number of incidences between points and curves, *Combinat. Probab. Comput.* 7 (1998), 121–127.

[33] N. D. Phuong, P. V. Thang, and L. A. Vinh, Incidences between points and generalized spheres over finite fields and related problems, in arXiv:1410.7899.

[34] M. Rudnev and J. M. Selig, On the use of Klein quadric for geometric incidence problems in two dimensions, in arXiv:1412.2909.

[35] M. Sharir, A. Sheffer, and J. Zahl, Improved bounds for incidences between points and circles, *Combinat. Probab. Comput.* 24 (2015), 490–520. Also in arXiv:1208.0053.

[36] M. Sharir and N. Solomon, Incidences between points and lines in three dimensions, *Proc. 31st Annu. Sympos. Comput. Geom.* (2015), 553–568. Also in arXiv 1501.02544.

[37] M. Sharir and N. Solomon, Incidences between points on a variety and planes in three dimensions, in arXiv 1603.04823.

[38] M. Sharir and J. Zahl, Cutting algebraic curves into pseudo-segments and applications, in preparation.

[39] J. Solymosi and V. H. Vu, Near optimal bounds for the Erdős distinct distances problem in high dimensions, *Combinatorica* 28 (2008), 113–125.

[40] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinat. Probab. Comput.* 6 (1997), 353–358.

[41] H. E. Warren, Lower bound for approximation by nonlinear manifolds, *Trans. Amer. Math. Soc.* 133 (1968), 167–178.

[42] J. Zahl, An improved bound for the number of point-surface incidences in three dimensions, *Contrib. Discrete Math.* 8(1) (2013), 100–121.